

SPATIO-TEMPORAL DEEP GRAPH INFOMAX

Felix L. Opolka*, **Aaron Solomon***, **Cătălina Cangea**, **Petar Veličković**, **Pietro Liò**

Department of Computer Science and Technology

University of Cambridge

{f1o23, acs207, ccc53, pv273, p1219}@cam.ac.uk

R Devon Hjelm

Microsoft Research

Mila – Québec Artificial Intelligence Institute

devon.hjelm@microsoft.com

ABSTRACT

Spatio-temporal graphs such as traffic networks or gene regulatory systems present challenges for the existing deep learning methods due to the complexity of structural changes over time. To address these issues, we introduce Spatio-Temporal Deep Graph Infomax (STDGI)—a fully unsupervised node representation learning approach based on mutual information maximization that exploits both the temporal and spatial dynamics of the graph. Our model tackles the challenging task of node-level regression by training embeddings to maximize the mutual information between patches of the graph, at any given time step, and between features of the central nodes of patches, in the future. We demonstrate through experiments and qualitative studies that the learned representations can successfully encode relevant information about the input graph and improve the predictive performance of spatio-temporal auto-regressive forecasting models.

1 INTRODUCTION AND RELATED WORK

Using deep learning techniques to analyze spatio-temporal data has shown promising results on classification and regression tasks, in contexts such as road and traffic networks (Shi & Yeung, 2018), gene expression data (Dutil et al., 2018), and Internet networking (Boutaba et al., 2018). Previous work has leveraged supervised learning techniques, largely by combining auto-regressive models with graph convolutional layers of different types (Yu et al., 2018; Li et al., 2018; Zhang et al., 2018). In this work, we address this task in the unsupervised learning setting and propose an approach for learning node representations for a spatio-temporal graph in a fully unsupervised fashion, encoding useful information that increases performance of a downstream forecasting model.

While neural network based learning methods for graph-structured data have received substantial attention, previous work has largely focused on supervised classification tasks. Veličković et al. (2019) have recently proposed Deep Graph Infomax (DGI)—an unsupervised representation learning approach for nodes in non-temporal graphs—and achieved state-of-the-art performance on classification benchmarks. Unlike previous methods (Mutlu & Oghaz, 2018), DGI does not rely on a random walk or adjacency-based methods and instead uses graph convolutions (Kipf & Welling, 2016) to build on the deep mutual information maximization principle described by Hjelm et al. (2019). So far, DGI has only been applied to non-temporal graphs in the node classification setting. In this work, we adapt the mutual information maximization principle to spatio-temporal graphs and show that the learned embeddings can encode valuable information for node regression tasks. We compare our model to a baseline auto-regressive model that only exploits temporal information and thus show that we can encode relevant spatial information in a fully unsupervised manner.

*Equal contribution




Figure 1: Visualization of the unsupervised training procedure. The embeddings $\mathbf{H}^{(t)}$ for the nodes at time step t are computed by the encoder \mathcal{E} . Given the embeddings and the uncorrupted raw features $\mathbf{X}^{(t+k)}$, the discriminator \mathcal{D} is trained to output that the sample is positive. When given the embeddings and the corrupted raw features $\tilde{\mathbf{X}}^{(t+k)}$, the discriminator is trained to output that the sample is negative.

2 BACKGROUND

2.1 PROBLEM SETTING

The node regression task takes the form of a prediction on a graph $G = (V, E, \mathbf{W})$ with node set V , edge set E , and weighted adjacency matrix \mathbf{W} . Node features change over time and hence features at time step t are given by $\mathbf{X}^{(t)}$. Given the features from the most recent T' time steps, the task is to predict the features of the next T time steps using a function $f(\cdot)$, potentially parameterized by a neural network:

$$[\mathbf{X}^{(t-T'+1)}, \dots, \mathbf{X}^{(t)}; G] \xrightarrow{f(\cdot)} [\mathbf{X}^{(t+1)}, \dots, \mathbf{X}^{(t+T)}]. \quad (1)$$

2.2 LEARNING REPRESENTATIONS VIA MUTUAL INFORMATION MAXIMIZATION

Deep InfoMax (DIM, Hjelm et al., 2019) is a recent approach for unsupervised representation learning that derives embeddings by maximizing the mutual information between the output of an encoder and local patches of the input. DIM builds on Mutual Information Neural Estimation (MINE, Belghazi et al., 2018), which formulates an estimate $\hat{I}(X; Y)$ for the mutual information between random variables X, Y using neural networks. These estimates are obtained by training a classifier (a.k.a. the *discriminator* or *statistics network*) to distinguish between samples from the joint distribution and the product of marginals. DIM applies this approach to representation learning by training both the encoder and the discriminator to maximize the mutual information between the random variables corresponding to local input patches and the embeddings. Deep Graph Infomax (DGI) extends this representation learning technique to non-temporal graphs, finding node embeddings that maximize the mutual information between local patches of the graph and summaries of the entire graph. Here, we build on these methods and propose a representation learning technique for spatio-temporal graphs. Furthermore, unlike in previous work, we evaluate our embeddings in the regression rather than classification setting.

3 SPATIO-TEMPORAL DEEP GRAPH INFOMAX

We extend the DGI approach by adapting it to spatio-temporal graphs and refer to our method as *spatio-temporal deep graph infomax* (STDGI). At each time step, representations are trained for each node in the graph in a fully unsupervised fashion. Similarly to DIM, we train the encoder to maximize the mutual information between patches in the graph at a particular time step t and the raw features of the same node at a future time step $t + k$. The goal is to aggregate, for each node, the information from its neighbourhood that is most relevant for predicting its features in the future.

3.1 ARCHITECTURE COMPONENTS

The unsupervised training setup is equivalent to the one in DIM or DGI. An encoder \mathcal{E} computes embeddings $\mathbf{h}_v^{(t)} \in \mathbb{R}^K$ for each node $v \in V$ at each time step t , using node features \mathbf{X} and the graph structure \mathbf{W} . The discriminator \mathcal{D} then receives pairs $(\mathbf{h}_v^{(t)}, \mathbf{x}_v^{(t+1)})$ containing the embedding and raw features of the same node at the current and next time step, respectively. We refer to such a pair of embedding and raw features as a *positive sample*, if both were drawn from the same graph. A *negative sample* will then consist of an embedding and raw features, where the latter are obtained from a corrupted version of the graph, derived by randomly permuting the node features of the graph at each time step. Positive samples can be understood as being drawn from the *joint distribution* of embeddings and raw features, whereas negative samples are drawn from the *marginal distributions*. The discriminator outputs a score corresponding to whether a given pair represents a positive or negative sample; both the encoder and discriminator are trained jointly to distinguish between positive and negative samples by minimizing the binary cross entropy loss. This maximizes the mutual information between the embeddings and the raw features of the next time step (Poole et al., 2018; Belghazi et al., 2018; Hjelm et al., 2019).

During supervised training, the embeddings $\mathbf{h}_v^{(t)}$ output by the encoder are concatenated with the raw features $\mathbf{x}_v^{(t)}$. In order to use the learned representations for our considered task, the resulting features $\mathbf{X}^{*(t)} = \mathbf{X}^{(t)} \oplus \mathbf{H}^{(t)}$ serve as input to a *downstream supervised regressor*.

4 EXPERIMENTS

We devise an evaluation setup for the *traffic forecasting task* to determine whether embeddings successfully encode spatial information of the graph that is relevant for making more accurate predictions. From a high-level point of view, the graph structure corresponds to a network of traffic sensors, where nodes are individual traffic sensors. An edge between two nodes is added when the distance between the two corresponding sensors is below a certain threshold. The time series of node features are given by the traffic measurements of each sensor over time.

4.1 EXPERIMENTAL SETUP

We use the METR-LA dataset (Jagadish et al., 2014), which contains data recorded by 207 traffic sensors. The traffic measurements were aggregated into five minute intervals and consist of traffic speed and the time of day. The graph is given by a directed, weighted adjacency matrix. The edge weights are the exponentially decaying distances along the roads. For more details on the data set, we refer to (Jagadish et al., 2014, E.1). Given the past 12 time steps (corresponding to measurements over 1h), the predictor has to forecast the traffic speeds at the next 12 time steps.

The encoder consists of a linear layer applied to each $\mathbf{x}_v^{(t)}$, followed by two graph convolutional layers applied to each $\mathbf{X}^{(t)}$. The discriminator is a two-layer fully-connected neural network, which concatenates the embedding and raw features of each pair and outputs whether the pair is a positive or negative sample. We chose to train three separate discriminators of the same architecture. Each discriminator compares the embedding to the raw feature of the same node k steps in the future where $k = 1, 3, 6$. For the downstream regressor, we employ an LSTM seq2seq model (Sutskever et al., 2014), which operates on the time series of each node in isolation. As a baseline, we compare our regressor to one with an identical configuration that receives as input *only the raw features* (rather than the concatenation of the raw features and embeddings). More details on the experimental setup can be found in Appendix A.

4.2 RESULTS

Results for the LSTM regressor that only uses raw features (baseline) and for the one that uses raw features concatenated with STDGI embeddings are shown in Table 1. We find that regressors using STDGI embeddings achieve lower predictions errors for all time horizons considered, and that the improvements *become more pronounced for larger time horizons*. This suggests that STDGI extracts embeddings with useful long-term features that improve upon the future predictive ability of raw data alone. All performances increases documented here are significant at $\alpha = 0.01$.

Method	15 min	30 min	60 min
MAE			
LSTM Baseline	3.67 ± 0.035	4.88 ± 0.017	6.53 ± 0.015
STDGI	3.61 ± 0.0038	4.80 ± 0.044	6.41 ± 0.016
RMSE			
LSTM Baseline	8.51 ± 0.037	10.66 ± 0.0076	13.04 ± 0.018
STDGI	8.28 ± 0.0042	10.42 ± 0.0055	12.79 ± 0.018
MAPE			
LSTM Baseline	$7.8 \pm 0.03\%$	$9.9 \pm 0.06\%$	$13.1 \pm 0.09\%$
STDGI	$7.6 \pm 0.01\%$	$9.5 \pm 0.01\%$	$12.5 \pm 0.04\%$

Table 1: Comparison of the LSTM baseline regressor and the STDGI method on METR-LA. We compute the MAE, RMSE, and MAPE of predictions for three different time horizons.



Figure 2: t-SNE visualization of embeddings from randomly selected time points spanning the available time series. Embeddings were visualized in two dimensions via t-SNE and colored by speeds at three timesteps (15 minutes) in the future to demonstrate the relationship between embedding similarity and true speed similarity.

These results are further supported by the qualitative study in Figure 2, which illustrates t-SNE-processed embeddings colored by future time point speeds. As indicated by their color, closely clustered embeddings generally share similar speeds. This suggests that embedding similarity is a proxy for speed similarity, and thus that embeddings are learning useful long-term information.

5 CONCLUSION

We have presented STDGI, an approach for learning embeddings of nodes in a graph that evolves over time, which leverages mutual information maximization and performs node regression in a spatio-temporal prediction context. We demonstrate that an auto-regressive seq2seq model operating on the time axis achieves higher predictive performance when making use of STDGI embeddings, thus confirming that the method successfully encodes valuable information over that provided by the standard baseline, even in the more challenging setting of regression. Moreover, the results show that STDGI is able to capture intrinsic and helpful properties of traffic flow by building increasingly stronger representations of the graph in relation to the baseline, as the prediction time horizon becomes larger. Our model represents both a generalization of DGI to spatio-temporal settings as well as a successful extension of the method to perform regression, in addition to classification. Future work will aim to find embeddings that further increase the accuracy of the downstream regressor and provide high-quality representations for multiple predictive tasks.

REFERENCES

- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeswar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. Mutual Information Neural Estimation. *arXiv preprint arXiv: 1801.04062*, 2018.
- Raouf Boutaba, Mohammad A. Salahuddin, Noura Limam, Sara Ayoubi, Nashid Shahriar, Felipe Estrada-Solano, and Oscar M. Caicedo. A comprehensive survey on machine learning for networking: evolution, applications and research opportunities. *Journal of Internet Services and Applications*, 9(1):16, 2018.
- Francis Dutil, Joseph Paul Cohen, Martin Weiss, Georgy Derevyanko, and Yoshua Bengio. Towards Gene Expression Convolutions using Gene Interaction Graphs. Technical report, 2018.
- R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *International Conference on Learning Representations*, 2019.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, 1997.
- H. V. Jagadish, Johannes Gehrke, Alexandros Labrinidis, Yannis Papakonstantinou, Jignesh M. Patel, Raghu Ramakrishnan, and Cyrus Shahabi. Big Data and Its Technical Challenges. *Commun. ACM*, 57(7):86–94, July 2014.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv: 1412.6980*, 2014.
- Thomas N Kipf and Max Welling. Semi-Supervised Classification with Graph Convolutional Networks. *arXiv preprint arXiv:1609.02907*, 2016.
- Yaguang Li, Rose Yu, Cyrus Shahabi, and Yan Liu. Diffusion Convolutional Recurrent Neural Network: Data-Driven Traffic Forecasting. In *International Conference on Learning Representations (ICLR '18)*, 2018.
- Ece C Mutlu and Toktam A Oghaz. Review on Graph Feature Learning and Feature Extraction Techniques for Link Prediction. 2018.
- Ben Poole, Sherjil Ozair, Aäron van den Oord, Alexander A Alemi, and George Tucker. On variational lower bounds of mutual information. *NeurIPS Workshop on Bayesian Deep Learning*, 2018.
- Xingjian Shi and Dit-Yan Yeung. Machine Learning for Spatiotemporal Sequence Forecasting: A Survey. *arXiv preprint arXiv: 1808.06865*, 2018.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pp. 3104–3112, 2014.
- Petar Veličković, William Fedus, William L. Hamilton, Pietro Liò, Yoshua Bengio, and R Devon Hjelm. Deep Graph Infomax. In *International Conference on Learning Representations*, 2019.
- Bing Yu, Haoteng Yin, and Zhanxing Zhu. Spatio-Temporal Graph Convolutional Networks: A Deep Learning Framework for Traffic Forecasting. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI-18*, pp. 3634–3640. International Joint Conferences on Artificial Intelligence Organization, 2018.
- Jiani Zhang, Xingjian Shi, Junyuan Xie, Hao Ma, Irwin King, and Dit-Yan Yeung. GaAN: Gated Attention Networks for Learning on Large and Spatiotemporal Graphs. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pp. 339–349, 2018.

A DETAILS ON EXPERIMENTAL SETUP

The METR-LA (Jagadish et al., 2014) dataset contains data recorded by 207 traffic sensors throughout Los Angeles County, from March 1st, 2012 to June 20th, 2012. In our experiments, we use the canonical split of the dataset into training, validation, and test set containing 23974, 3425, and 6850 samples, respectively.

All layers in the encoder contain 64 hidden units and the embeddings size is 128. The two fully-connected layers of the discriminator contain 6 and 1 hidden units, respectively. The seq2seq downstream regressor consists of a single LSTM layer (Hochreiter & Schmidhuber, 1997) with 64 hidden units.

All models are trained with a batch size of 64 for 120 epochs, using the Adam optimizer (Kingma & Ba, 2014). The unsupervised training of embeddings is carried out for 100 epochs, with an initial learning rate of $1e^{-3}$ that is reduced by a factor of $\frac{1}{10}$ every 30 epochs after the first 20. The supervised models use the mean absolute error (MAE) over the entire horizon of 12 steps as the loss function. The learning rate is initially $1e^{-2}$ and decreases by a factor of $\frac{1}{10}$ every 30 epochs after the first 20.